


Introduction to the Accelerated Strategic Computing Initiative Academic Strategic Alliance Program

Charles McMillan

This paper was prepared for submittal to the
American Institute of Aeronautics and Astronautics
Norfolk, VA
June 28 - July 1, 1999

April 27, 1999



This is a preprint of a paper intended for publication in a journal or proceedings.
Since changes may be made before publication, this preprint is made available with
the understanding that it will not be cited or reproduced without the permission of the
author.

DISCLAIMER

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial products, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

INTRODUCTION TO THE ACCELERATED STRATEGIC COMPUTING INITIATIVE ACADEMIC STRATEGIC ALLIANCE PROGRAM

Charles McMillan
Associate B Division Leader for Computational Physics
Lawrence Livermore National Laboratory*
Livermore, CA 94550

ABSTRACT

The Accelerated Strategic Computing Initiative of the Department of Energy has chosen to fund major simulation projects at five universities as one of its key strategies. These projects were selected through a competitive process in the spring of 1997 with initial funding to the projects beginning in the autumn of 1997.

INTRODUCTION

At its inception, the Accelerated Strategic Computing Initiative (ASCI) chose five strategies as organizing principles.

- Focus on 3-D full physics, full systems applications
- Focus on the high-end of computing
- Use, leverage and adapt technologies to create problem-solving environments
- Encourage openness
- Execute ASCI as a 3-Lab, single-program activity

The Academic Strategic Alliance Program (ASAP) is the product of two of these strategies, "Encourage openness" and "Focus on 3-D full physics, full systems applications". The importance of engaging the best minds to work on problems relevant to ASCI was recognized as a key to the future success of ASCI. Furthermore, while many of these minds reside at the DOE Defense Programs Labs, it is clearly the case that many more are in the university community. Hence, the focus on "openness" as a strategy.

The nature of the problems faced by the DOE Labs drove the emphasis on "3-D full physics, full system applications". As historical approaches based on experiment have become unavailable to the Labs, simulation has moved to bridge the gap. However, it is no longer adequate to simulate the physics piecemeal, work in reduced dimensionality, or deal only with system components. Rather, it is now necessary to simulate the integrated system in much greater detail. The full-system nature of these problems guided the selection process in ASAP.

The goals of the alliance program grew from these two ASCI goals.

- Establish within the academic community the practices of large-scale modeling and simulation
- Accelerate advances in basic sciences, mathematics, and computer science for computational science and engineering
- Establish technical coupling of ASCI projects at the Lab with Strategic Alliance partners
- Strengthen training and research in areas of interest for Defense Programs.

In many ways, ASAP is a component in the answer to the types of questions asked by thoughtful reviewers of the ASCI program, "You are developing computational models for nuclear weapons simulations that you can never test. Can you develop comparable models for unclassified systems that can be tested and validated?" At \$15-20 million for each project spread over five years, these alliances are the basis for a new relationship between Universities and the labs as we answer this question.

* This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract No. W-7405-Eng-48.

ALLIANCES

The Department of Energy, working with the Defense Programs Labs, Lawrence Livermore National Lab, Los Alamos National Lab, and Sandia National Lab, began the process leading to the Alliance Program in late 1996 with a pre-proposal conference inviting university participation. In addition, this conference gave potential alliance partners briefings on ASCI and examples of the types of projects that were expected.

Pre-proposals were solicited from the university community that resulted in 48 high quality proposals. A first round of feedback on the pre-proposals went to the university community from reviewers in January, 1997 in preparation for the actual proposal process. Based on this feedback, 21 proposals were submitted for review during April, 1997. From these, 7 were selected for site reviews culminating in the announcement of the selection of 5 alliance partners in late July, 1997. Funding for these projects began in the fall of 1997.

The key factors in the selection process were

- Technical quality of the proposal and the PIs
- Presence of an overarching, multi-disciplinary problem requiring high performance computing
- Relevance to ASCI.

The first of these requires no comment since it is common to all scientific competitions. The other two, however, merit comment.

The presence of an overarching problem was evaluated by examining the complexity of the system that was proposed for study. Proposals that would pull together multiple academic disciplines, for example, physics, chemistry, engineering, and computer science, were favored over those that tended to be very narrow, for example hydrodynamics. The purpose of this portion of the selection process was to encourage the academic community to tackle problems of comparable difficulty to those faced by the Labs. This was viewed as being particularly important as a training ground for students. In addition, it was essential that these projects *require* high performance parallel computing for their success.

The requirement for projects that were "ASCI relevant" favored projects that sampled physical, chemical, and engineering problems that intersected the domains most important for the success of ASCI at the Labs. Here, the expectation of relevance was not only for training students, but also for the exchange of ideas and algorithms between the Universities and the Labs.

SELECTIONS

The peer review process provided selectable projects from which DOE selected five based on programmatic requirements. The five are:

University of Chicago: The simulation of astrophysical thermonuclear flashes

University of Illinois at Urban/Champaign: The simulation of advanced solid rocket motors in both normal and abnormal operation

University of Utah: The simulation of accidental fires and explosions focussing on fires involving energetic materials and structures

California Institute of Technology: The simulation of the dynamic response of materials focussing on explosively driven shocks

Stanford: Simulation of gas turbine engines including complex fluid flows, turbulence and chemistry.

Representative papers from each of these centers are included in this section of the proceedings that provide details of the activities at each of these centers.

To assist in ensuring the success of ASAP, DOE and the Labs have assembled Technical Support Teams (TST), one for each university, that meet with the university alliance teams on a regular basis. The TST not only tracks the progress of the projects, but is also responsible for acting as the interface between the university community and the Labs, making each aware of the work being done by the other. Based on the reports of these teams, all of the projects have made a good start during the first year of funding.